Bayesian statistics

1. (Laplace, 1786) Considering male and female births in Paris, Laplace wants to test whether the probability p of a male birth is above $1/2$. For 251527 male and 241945 female births, he assigns p a uniform prior distribution on $[0, 1]$.

   (a) Derive the posterior distribution of p.

   (b) Compute the posterior probability that $p \leq 1/2$.

2. Let X be $\mathsf{N}(\theta, \sigma^2)$-distributed, and suppose that the prior distribution on θ is $\mathsf{N}(\mu, \tau^2)$. Here we assume that μ, σ^2, and τ^2 are all known. Find the posterior distribution of θ and compute Bayes estimator $\hat{\theta} = \mathbb{E}(\theta \mid X)$ as well as $\mathbb{V}(\theta \mid X)$.

3. Let $X = (X_1, X_2, \ldots, X_n)$ be a vector of independent *Pareto-distributed* random variables with common density
$$p_\theta(x) = \frac{1}{\theta x^{(1+1/\theta)}}, \quad x > 1,$$
where $\theta > 0$ is an unknown parameter. Consider a Bayesian framework where θ has a prior distribution with density
$$\pi(\theta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{-\alpha-1} \exp\left(-\frac{\beta}{\theta}\right), \quad \theta > 0 ,$$
with $\alpha > 1$ and $\beta > 0$ being fixed hyperparameters; this means that θ has *inverse gamma distribution* with mean $\beta/(\alpha - 1)$ (code: $\mathsf{IG}(\alpha, \beta)$). Compute Bayes estimator of θ.

MCMC for Bayesian computation

4. We consider a *variance components model* comprising

   – fixed constants $\mu_0 \in \mathbb{R}$ and $(\alpha_1, \beta_1, \alpha_2, \beta_2, \sigma_0^2) \in \mathbb{R}_+^*$,

   – three hyperparameters
$$\sigma_\theta^2 \sim \mathsf{IG}(\alpha_1, \beta_1), \quad \sigma_\varepsilon^2 \sim \mathsf{IG}(\alpha_2, \beta_2), \quad \mu \sim \mathsf{N}(\mu_0, \sigma_0^2),$$

   – I further components $(\theta_i)_{i=1}^I$ being conditionally independent given the above hyperparameters with
$$\theta_i \sim \mathsf{N}(\mu, \sigma_\theta^2),$$

   – data $\mathbf{Y} = (Y_{ij})_{i=1,\ldots I; j=1,\ldots,J}$ which are assumed to be distributed as
$$Y_{ij} \sim \mathsf{N}(\theta_i, \sigma_\varepsilon^2)$$
   conditionally independently given the parameters.

   A graphical representation of the model is as follows:
$$\begin{array}{ccccc} & & \theta_1 & \to & (Y_{1,j})_{j=1}^J \\ & \nearrow & \vdots & & \vdots \\ \mu & \to & \theta_i & \to & (Y_{i,j})_{j=1}^J \\ & \searrow & \vdots & & \vdots \\ & & \theta_I & \to & (Y_{I,j})_{j=1}^J \end{array}$$

We are interested in the posterior

$$f(\sigma_\theta^2, \sigma_\varepsilon^2, \mu, \theta_1, \ldots, \theta_I \mid \mathbf{Y}),$$

and aim at sampling from this complex distribution using a Gibbs sampler. For this purpose, we proceed as follows.

(a) Write $f(\sigma_\theta^2, \sigma_\varepsilon^2, \mu, \theta_1, \ldots, \theta_I \mid \mathbf{Y})$ up to a normalizing constant.

(b) Use the expression in (a) for deriving the full conditionals

(i) $f(\mu \mid \sigma_\theta^2, \sigma_\varepsilon^2, \theta_1, \ldots, \theta_I, \mathbf{Y})$,

(ii) $f(\sigma_\theta^2 \mid \sigma_\varepsilon^2, \mu, \theta_1, \ldots, \theta_I, \mathbf{Y})$,

(iii) $f(\sigma_\varepsilon^2 \mid \sigma_\theta^2, \mu, \theta_1, \ldots, \theta_I, \mathbf{Y})$,

(iv) $f(\theta_i \mid \sigma_\varepsilon^2, \sigma_\theta^2, \mu, \theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I, \mathbf{Y})$, for $i = 1, \ldots, I$.

(c) Write a pseudo-code describing one full sweep (i.e., one iteration) of the resulting Gibbs sampler. If you have time, implement the resulting algorithm in MATLAB for some data record $\mathbf{Y}$ obtained through simulation under some suitably chosen values of $(\mu_0, \alpha_1, \beta_1, \alpha_2, \beta_2, \sigma_0^2)$. Illustrate the marginal posteriors using normalized histograms. Do the marginals cluster around the true parameters?